UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/814,308	03/22/2001	Alan Paul Rolleston Phillips	ASTB-0044	4836

7590 10/20/2006

Woodcock Washburn Kurtz
Mackiewicz & Norris LLP
46th Floor
One Liberty Place
Philadelphia, PA 19103

EXAMINER	
VAN DOREN, BETH	
ART UNIT	**PAPER NUMBER**
3623	

DATE MAILED: 10/20/2006

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary	Application No.	Applicant(s)	
	09/814,308	PHILLIPS, ALAN PAUL ROLLESTON	
	Examiner	Art Unit	
	Beth Van Doren	3623	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *27 July 2006*.

2a)☒ This action is **FINAL**. 2b)☐ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☒ Claim(s) *18-34* is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) *18-34* is/are rejected.

7)☐ Claim(s) _____ is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☐ The specification is objected to by the Examiner.

10)☐ The drawing(s) filed on _____ is/are: a)☐ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11)☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☐ All b)☐ Some * c)☐ None of:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☐ Notice of References Cited (PTO-892)
2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
3) ☐ Information Disclosure Statement(s) (PTO/SB/08) Paper No(s)/Mail Date _____.
4) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____.
5) ☐ Notice of Informal Patent Application
6) ☐ Other: _____.

DETAILED ACTION

1. The following is a final office action in response to the communications received 07/27/2006. Claims 18, 24-26, 28-29, and 31-33 have been amended. Claim 34 has been added. Claims 18-34 are now pending in this application.

Response to Amendment

2. Applicant's amendment to the specification is sufficient to overcome the specification objections set forth in the previous office action.

3. Applicant's amendments to claims 31-32 are not sufficient to overcome the claim objections set forth in the previous office action. These objections have been updated based on the amendments and reasserted below.

4. Applicant's amendments to claims 18 and 33 are sufficient to overcome the 35 USC § 112, second paragraph, rejections of claims 18-26, 29-30, and 33 set forth in the previous office action.

5. Applicant's response regarding claim 27 (which includes discussion of the amendment made to claim 18) is not sufficient to overcome the 35 USC § 112, second paragraph, rejection of claim 27 set forth in the previous office action. Even with the amendments made to claim 18, it is still not clear as to how the inclusion of a robot relates to the recited method steps in the body of claim 18, since it is not clear how the specifics of the structure of the robot are affected by the method recited. Therefore, the 35 USC § 112, second paragraph, rejection is maintained below.

6. Applicant's amendments to claims 28, 29, 31, and 32 are not sufficient to overcome the 35 USC § 112, second paragraph, rejections set forth in the previous office action. Therefore, these rejections have been updated based on the amendments and reasserted below.

7. Applicant's amendments to claim 18 are sufficient to overcome the 35 USC § 101 rejections of claims 18-30 and 33 set forth in the previous office action.

8. Since Examiner is construing the system of claim 31 as the same system that runs the method of claim 18 (asserted below in the 35 USC § 112, second paragraph, rejection of claim 31), Applicant's amendments to claim 31 are sufficient to overcome the 35 USC § 101 rejections of claims 31-32 set forth in the previous office action. However, Applicant's response to the 35 USC § 112, second paragraph, rejections and the claim objection below may cause examiner to revisit such rejections.

Claim Objections

9. Claims 31-32 are objected to under 37 CFR 1.75(c), as being of improper dependent form for failing to further limit the subject matter of a previous claim. Applicant is required to cancel the claims or amend the claims to place the claims in proper dependent form, or rewrite the claims in independent form.

As per claim 31, claim 31 recites "a system having a control apparatus that is programmed to control the objective function of the system according to claim 18". From this language, claim 31 is not required to include every limitation of parent claim 18, since the system of claim 31 only appears to be a system that controls the objective function of claim 18, and thus does require the other method limitations of claim 18, such as performing the monitoring step or the storing step. Further, while the method of claim 18 is required to perform the actions of the claim, claim 31 is directed to the structure of the system, and thus the acts are not necessarily performed. Thus, claim 31 may be infringed without claim 18 also being infringed.

Claim 32 depends from claim 31, and recites that "the system comprises a robot". Therefore, the robot of claim 32 is included in the system of claim 31 and does not remedy claim 31's deficiencies, as set forth above, because the robot does not specifically include and perform every limitation of parent claim 18, and is merely included in a system the controls the objective function of claim 18. Thus, claim 32 may be infringed without claim 18 also being infringed.

Claim Rejections - 35 USC § 112

10. The following is a quotation of the second paragraph of 35 U.S.C. 112:

> The specification shall conclude with one or more claims particularly pointing out and distinctly claiming the subject matter which the applicant regards as his invention.

11. Claims 27-29 and 31-32 are rejected under 35 U.S.C. 112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which applicant regards as the invention.

Claim 27 recites "a method of controlling a system wherein the system comprises a robot". It is still unclear as to how the inclusion of a robot relates to the recited method steps in the body of claim 18. Therefore, clarification is required. For examination purposes, it has been construed that the output of the method of claim 18 controls an external robot.

Claim 28 recites systems having ranks of control arranged in hierarchies, wherein candidate actions "at the lowest level in the hierarchy represents the output candidate action selected to be performed by the system, and wherein the candidate action of a rank of control not at the lowest level in the hierarchy represents the selection of a lower rank of control in the hierarchy". It is unclear as to what specifically is occurring in this claim. In claim 28, there appears to be at least two ranks of control (one at a lowest level and one at a level not at the lowest level). The claim states that each of these ranks of control is performed according to the

method of claim 18. However, claim 28 also states that "a candidate action of a rank of control at the lowest level [...] represents the output candidate action selected to be performed" and "a candidate action not at the lowest level [...] represents the selection of a lower rank of control". Therefore, it is unclear how if the same methodology from claim 18 is used to control both ranks of control, how these ranks of control would arrive at different output (the lowest level selects an action and the non-lowest level selects a rank of control). Therefore, it is unclear as to how claim 28 would be specifically implemented and how each individual rank of control would operate according to the method of claim 18. Clarification is required.

Claim 29 depends for claim 28, and is therefore rejected using the same rationale set forth above. Claim 29 further recites that "the monitored response performances of the lowest level ranks of control are all visible and accessible to the rank of control immediately above in the hierarchy, for the purposes of appraising the probability distribution of the response performance of all of said plurality of candidate actions". First, it is not specifically clear as to what rank of control the language "the rank of control immediately above" is referring, as claim 28 recites ranks of control at "the lowest level" and "not at the lowest level", but does not expressly state that the ranks of control "not at the lowest level" are immediately above "the lowest level". Therefore, it is unclear if the claim is referring to the ranks of control "not at the lowest level" or another rank of control by this limitation. Second, claim 29 recited "the monitored response performances", which refers back to claim 18. As discussed above, it is not specifically clear as to how the "ranks of control" concept functionally fits with the limitations of claim 18. Clarification is required.

Claim 31 recites a system having a control apparatus that is programmed to control the objective function of the system. It is clear in this claim limitation that the system in the language "control the objective function of the system" refers to the system recited in claim 18. However, it is not clear as to whether the system in the language "a system having a control apparatus" is the same system as the system of claim 18 or a second system that controls the method and system of claim 18. Clarification is required. For examination purposes, it has been construed that the system is the same system as that which runs the method of claim 18, with the control apparatus included in the system of claim 18.

Claim 32 recites that "a system according to claim 31 [...] comprises a robot". Since neither claim 18 nor claim 31 recite any language concerning a robot, it is unclear as to how the robot specifically affects the elements of claims 18 and/or 32. It is unclear as to whether the controlling of the objective function (claim 31) is implemented on the robot, if the robot is separate and merely controlled by the controlling of the objective function of claim 31, what the robot does to its surroundings, etc. Further, it is not specifically clear as to how the structure of the robot functionally interacts with the system of claim 31 and further the method of claim 18. Clarification is required. For examination purposes, it has been construed that the robot is within the system of claim 32 and controlled by the controlling of the objective function.

Because claims 28-29 are so indefinite, no art rejection is warranted as substantial guesswork would be involved in determining the scope and content of these claims. See In re Steele, 305 F.2d 859, 134 USPQ 292 (CCPA 1962); Ex parte Brummer, 12 USPQ 2d, 1653, 1655 (BdPatApp&Int 1989); and also In re Wilson, 424 F.2d 1382, 165 USPQ 494 (CCPA

1970). Prior art pertinent to the disclosed invention is nevertheless cited and applicants are reminded they must consider all cited art under Rule 111(c) when amending the claims to conform with 35 U.S.C. 112.

Response to Arguments

12. Applicant's arguments with regards to the 35 USC § 103 rejections based on Merriman et al. (U.S. 2002/0099600) in view of Eppen et al. (*Quantitative Concepts for Management*) have been fully considered, but they are not persuasive. In the remarks, Applicant argues that (1) Merriman et al. in view of Eppen et al. do not teach and suggest the requirements of learning efficiency or its role in the overall performance of a self-regulating decision system, (2) Merriman et al. in view of Eppen et al. does not teach and suggest the claimed minimization in the growth of regret and hence do not teach all the features of the claimed invention, (3) Eppen et al. does not teach or suggest optimizing an objective function by assessing the probability distributions of all the candidate actions in order to control the growth in regret, "where regret is a term [...] candidate action" (see last element, claim 1) or the merit of actually taking an action which is expected to offer a lower immediate payoff because the value of new information gained exceeds the loss resulting from deliberately taking an action with a lower pay-off (i.e. an exploration-exploitation tradeoff that is specially incorporated in regret), (4) Eppen et al. assumes that probabilistic information about the state of nature is obtained by means other than the design process itself, which is unlike the claimed system where no external information is available and all information used is collected by the system from observations, (5) Eppen et al. does not disclose that the payoff of regret is not known, (6) Examiner has not established a prima facie case of obviousness since Merriman et al. in view of Eppen et al. does not teach the

minimization in the growth of regret (and thus does not teach each and every limitation) and further does not provide a motivation to combine.

In response to argument (1) that the references fail to show certain features of applicant's invention, it is noted that the features upon which applicant relies (i.e., learning efficiency, self-regulating decision system) are not recited in the rejected claims. Although the claims are interpreted in light of the specification, limitations from the specification are not read into the claims. See *In re Van Geuns*, 988 F.2d 1181, 26 USPQ2d 1057 (Fed. Cir. 1993).

Claims 18 and 33 do not specifically recite the terms "learning efficiency" and "self-regulating". Claims 18 and 33 recite steps of monitoring response performance, storing a representation of response performance, choosing a next candidate action to perform based on a probability distribution of the response performance of all the candidate actions and an objective function, performing the chosen candidate action, and repeating. While these steps are implemented using a control apparatus and a method, and while the actions are performed by a system, the claims contain no specific recitation requiring them to be completely self-regulated by the computer system (with no human intervention). Further, while the steps above are iterative and a next candidate action to be performed is chosen based on a probability distribution of the response performance, there is no specific recitation of "learning efficiency" or how the method and apparatus specifically learn from the iterative cycle, beyond the system storing response performance and using it somehow in the choosing of a next action.

Merriman et al. discloses that the response performance to a candidate action (advertising) is monitored and stored in the historical database of the system. See paragraphs 0010, 0015-6, 0031, 0033-4. Based on the knowledge gained and stored concerning response

performance, a next action (ad) is chosen to be performed by the system to optimize an objective function by assessing, using a predictive model, empirical data to determine which action will maximize feedback/minimize economic loss after the chosen candidate action is performed based on historical response performances to date by the system. See paragraphs 0008, 0017-8, 0033, 0039, 0041-2. This is an iterative process, where the model is refined over time. Thus, Merriman et al. discloses these aspects, as claimed, by disclosing a system that uses an iterative process to refine a model over time (i.e. learning efficiency), this process performed via a system (decision system).

In response to argument (2), Examiner respectfully disagrees. Examiner points out that she did not rely on Merriman et al. to disclose regret (as explicitly stated in the 35 USC 103 rejections below). Eppen et al. was relied upon to disclose the concept of regret as well as to disclose the minimization (or lowest expected) growth in regret after the chosen candidate action is performed. See specifically page 511, section 1, which specifically states that when the decision maker/software knows the probability distribution on the state of nature, regret can be minimized. See also page 512-513. Therefore, Eppen et al. does expressly disclose the minimization of the growth/development of regret. The concept of regret with regards to the Eppen et al. reference and its combination with Merriman et al. will be further addressed in the subsequent arguments.

In response to argument (3), Examiner respectfully disagrees. Examiner first points out that neither claim 18 nor claim 33 recite limitations or language that cause the method, system, and/or apparatus to choose to perform an action which is expected to offer a lower immediate payoff because the value of new information gained exceeds the loss resulting from taking an

action with a lower pay-off. In fact, Examiner believes that, as currently recited, the claims require just the opposite. Element c) of claims 18 and 33 specifically recites "choosing which of the plurality of candidate actions is next performed so as to optimize said objective function by assessing [...] which candidate action is estimated to ***result in the lowest expected growth in regret after the chosen action is performed***". Therefore, the claims specifically require that after the chosen action is performed, the growth in regret is lowered. Since the claims contain no recitation of time period or long term effects, actively choosing an action that is expected to offer a lower immediate payoff would not satisfy the recited claim limitation above (that the action is estimated to lower the regret).

The last limitation of claims 18 and 33 recite "where ***regret*** is a term that represents a system performance measure that ***considers*** the relative ***merit of exploration of*** one or more apparently ***non-best candidate actions with respect to*** the relative ***merit of exploiting*** what appears to be the ***current best candidate action***". First, while this limitation discusses that regret represents ***a measure of consideration*** between exploring a non-best action relative to exploiting a current best action, this limitation does not recite that an active choice is made based on the consideration. Further, this limitation is not actively linked to the preceding steps of the claim, since elements a)-e) do not recite the use of a system performance measure or a choice of action being made based on such a measure or consideration. Finally, while this final limitation does include the consideration of exploration and exploitation of non-best versus best actions, this limitation never states that a system or decision maker would actually choose the exploration option, even if the limitation was actually performed. The limitation merely states that its existence is considered with respect to the exploitation of the best action. Based on the

discussion above concerning the claimed invention seeking to lower the growth in regret, there is no reason in the claims to believe that the exploration option would in fact be chosen, even if this limitation is active.

Looking at the Eppen et al. reference, which discusses well-known regret theory utilized in economic and decision theory, each possible decision (or action) has associated states of nature (outcomes). Therefore, in each set of possible decisions (or actions), an action is associated with an apparently best outcome and another action is associated with an apparently non-best outcome. They are the apparent best and non-best outcomes because they have not yet occurred, but it is evident that they will occur in this way. When generating a regret table, as shown and discussed on pages 510-511, regret is expressed as a performance measure, the table showing the merit (i.e. value, advantage, worth) of exploring a non-best action (decision) versus the merit of exploiting a best action (decision), as represented by the numbers in the table that reflect opportunity cost/loss. Therefore, Eppen et al. does disclose regret, as recited in the final limitation of claims 18 and 33.

Finally, Eppen et al. does disclose optimizing an objective function by assessing, using the probability distribution of the response performance of all of said plurality of candidate actions, which candidate action is estimated to result in the lowest expected growth in regret after the chosen candidate action is performed. Examiner notes that no specific objective function is recited in the claims; rather its functionality is defined, such that it is optimized to show which action results in the lowest growth in expected regret. See page 503, page 504, section 1, 511, section 1, wherein when the decision maker/software knows the probability distribution on the

state of nature, expected growth in regret can be minimized as reflected in a value function. See also page 512-513.

In response to argument (4), Examiner points out that Merriman et al. was relied upon for the steps of monitoring and storing response performance with respect to a candidate action, wherein historical data about response performance is stored in a historical database. Based on this data, Merriman et al. discloses updating the predictive model based on the feedback so as to choose a next candidate action that optimizes the objectives, using empirical data to determine which action will minimize economic loss after the actions occurrence. This assessment is based on historical response performances to date by the system. This process is repeated, refining the model over time based on observed on response performance. As discussed in paragraph 0042 of Merriman et al., the decision of action of the predictive model is based on historical statistical performance of actions and the probability of a positive response to the action to determine the expected return of the action. Thus, Merriman et al. was relied upon to teach the obtaining of data over time, where the data is based on internal system knowledge. Eppen et al. discloses the general process of using regret theory to make a decision, the process using probability distributions (expected responses) and expected returns based on specific actions' occurences. Thus, the knowledge gained during the iterative process of Merriman et al. (i.e. statistical performance of actions and probability of positive response) would be represented using the regret decision model and would allow the system of Merriman et al. to make a decision that minimizes the expected growth in regret. Examiner further points out that Eppen et al. does not expressly limit his teachings to a specific source of data.

In response to argument (5), Examiner respectfully disagrees. First, it is noted that the features upon which applicant relies (i.e., that the payoff of regret is not known) is not recited in the rejected claims. Therefore, Examiner is not specifically clear as to which claim limitation the Applicant is referring. Second, Eppen et al. discloses the use of regret (or opportunity cost/lost) in the consideration of what action to take with respect to a group of actions, wherein the value associated with regret is based on probabilities and an expectation of an outcome. Therefore, what specific payout of regret occurs is not actually known to the system.

In response to argument (6), Examiner respectfully disagrees. Examiner has provided art that teaches each and every limitation of the claimed invention, as explained above and as set forth below. Further examiner has provided motivation to combine the references, the motivation found within the references themselves. See paragraphs 0002, 0008, and 0010 of Merriman et al., which disclose which disclose the opportunity cost of poorly performing served advertisements (actions) and how the system of Merriman et al. allows for more efficient use of actions by monitoring results and using a predictive model. See page 503 of Eppen et al., which discloses a framework for analyzing a wide variety of management problems using available information about the problem and a measure of goodness of a selected action, providing a pragmatic and practical aid in decision making. Eppen et al. specifically equates regret with opportunity costs. Therefore, it would have been obvious to one of ordinary skill in the art at the time of the invention to use probability distributions and the theory of regret in the iterative predictive model of Merriman et al. in order to increase the efficiency of utilizing advertising/action space by providing a decision framework with which to analyze the various options.

13. Applicant's arguments with regards to the 35 USC § 103 rejections based on Merriman et al. in view of Eppen et al. and in further view of McClave et al. (*A First Course in Business Statistics*) have been fully considered, but they are not persuasive. In the remarks, Applicant argues that (7) McClave et al. does not discuss the use of the Student t distribution to regulate the exploration/exploitation tradeoff to deliver the lowest expected growth in regret.

In response to argument (7), McClave et al. was not relied upon to disclose the exploration/exploitation tradeoff to deliver the lowest expected growth in regret. Examiner notes, as discussed above with respect to argument (3) that this feature is not specifically recited in the pending claims in such a manner. McClave et al. was relied upon to teach distributions of populations using a Student's distribution with Student's t parameters, as set forth below.

14. Applicant's arguments with regards to the 35 USC § 103 rejections based on Merriman et al. in view of Eppen et al. and in further view of Jameson (U.S. 6,032,123) have been fully considered, but they are not persuasive. In the remarks, Applicant argues that (8) Jameson does not teach or suggest making decisions that trade-off the value of acquiring new information with the potential losses realized by ignoring other candidate actions.

In response to argument (8), Examiner points out that Jameson was not relied upon to disclose the trade-off in value of acquiring new information with the potential losses realized by ignoring other candidate actions. And, as discussed above with respect to argument (3), this feature is not specifically recited in the pending claims in such a manner. Jameson was relied upon to disclose the use of a Monte Carlo algorithm to provide understanding using "what if"

simulation to facilitate analysis of the response performance of candidate actions, as set forth below. See response to argument (3) above, which discusses the values of actions with respect to regret. Examiner notes the alternative language used in claim 23.

15. Applicant's arguments with regards to the 35 USC § 103 rejections based on Merriman et al. in view of Eppen et al. and in further view of Strickland et al. (U.S. 5,790,407) have been fully considered, but they are not persuasive. In the remarks, Applicant argues that (9) Strickland et al. does not discuss using regulation of the exploration/exploitation tradeoff to affect explicit control in the growth of regret.

In response to argument (9), Examiner points out that Strickland et al. was not relied upon to disclose the exploration/exploitation tradeoff. And, as discussed above with respect to argument (3), this feature is not specifically recited in the pending claims in such a manner. Strickland was relied upon to teach control systems for controlling external devices, such as robots, by comparing the response profile of the device to the actual response of the device, as set forth below. See response to argument (3) above, which discusses the trade-off in value with respect to the claims.

Claim Rejections - 35 USC § 103

16. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

17. Claims 18-21, 24-25, 30, 31, and 33-34 are rejected under 35 U.S.C. 103(a) as being unpatentable over Merriman et al. (U.S. 2002/0099600) in view of Eppen et al. (*Quantitative Concepts for Management*).

As per claim 18, Merriman et al. teaches a method of controlling a system to optimize an objective function thereof, the system being capable of performing a plurality of candidate actions and being capable of monitoring response performances of a performance of a respective candidate action, the method comprising the steps of:

a) monitoring response performance of a respective candidate action that is chosen to be performed by the system (See paragraphs 0010, 0015-6, 0033-4, wherein response to the action (direct advertising) is monitored);

b) storing, according to candidate action performed by the system, a representation of said monitored response performance (See paragraph 0031, 0033, wherein historical data about the response to an action is stored in the historical database of the system);

c) choosing which of the plurality of candidate actions is next performed by the system so as to optimize said objective function by assessing, using a predictive model, empirical data to determine which action will maximize feedback/minimize economic loss after the chosen candidate action is performed based on historical response performances to date by the system (See paragraphs 0008, 0017-8, 0033, 0039, 0041-2, wherein an action is chosen based on the current known performance);

d) commanding the system to perform the candidate action identified to be the next performed in step c) (See paragraphs 0008, 0017-8, 0033, 0039, 0041-2, wherein an action is chosen and performed based on the current known performance);

e) repeating steps a) to d) to control the system so as to substantially optimize the objective function of the system (See paragraph 0019, wherein the steps are iteratively repeated).

However, Merriman et al. does not expressly disclose optimizing said objective function by assessing, using the probability distribution of the response performance of all of said plurality of candidate actions, which candidate action is estimated to result in the lowest expected growth in regret after the chosen candidate action is performed, where regret is a term that represents a system performance measure that considers the relative merit of exploration of one or more apparently non-best candidate actions with respect to the relative merit of exploiting what appears to be the current best candidate action.

Eppen et al. discloses optimizing said objective function by assessing, using the probability distribution of the response performance of all of said plurality of candidate actions, which candidate action is estimated to result in the lowest expected growth in regret after the chosen candidate action is performed (See page 503, page 504, section 1, 511, section 1, wherein when the decision maker/software knows the probability distribution on the state of nature, regret could be minimized. See also page 512-513),

where regret is a term that represents a system performance measure that considers the relative merit of exploration of one or more apparently non-best candidate actions with respect to the relative merit of exploiting what appears to be the current best candidate action (See page 510-511, wherein regret represents the value of one of the non-best actions with respect to the value of the current best action (i.e. regret is the opportunity cost of not making the best decision)).

Merriman et al. teaches a method and apparatus that considers past performance data when automatically determining the next action to take. Merriman et al. uses a predictive model with which to make a decision, the predictive model using past performance information to deliver optimal actions, thus maximizes utilization of the actions. Eppen et al. discloses the use of regret (or opportunity cost/lost) in the consideration of what action to take with respect to a group of actions based on a set of conditions. It would have been obvious to one of ordinary skill in the art at the time of the invention to use probability distributions and the theory of regret in the iterative predictive model of Merriman et al. in order to increase the efficiency of utilizing advertising/action space by providing a decision framework with which to analyze the various options. See paragraphs 0002, 0008, and 0010 of Merriman et al. and page 503 of Eppen et al.

As per **claims 19-21**, Merriman et al. discloses c) choosing which of the plurality of candidate actions is next performed so as to optimize said objective function by assessing, using a predictive model, empirical data to determine which action will maximize feedback/minimize economic loss after the chosen candidate action is performed based on historical response performances to date (See paragraphs 0008, 0017-8, 0033, 0039, 0041-2, wherein an action is chosen based on the current known performance). However, Merriman et al. does not expressly disclose and Eppen et al. discloses:

As per claim 19, that c) includes assessing which candidate action is likely to result in the lowest expected growth in regret on the basis of a true best candidate action which has the mean of said probability distribution (See page 510-511 and 512-513, wherein regret is assessed to determine which action will result in the lowest regret using a probability distribution and expected values of regret);

As per claim 20, that step c) includes evaluating the cost or losses associated with presenting a lower performing candidate action and the gain or benefit associated with knowing the true position of the current best observed candidate action on said probability distribution (See page 510-511, wherein regret represents the value of one of the non-best actions with respect to the value of the current best action (i.e. regret is the opportunity cost of not making the best decision)).

As per claim 21, that step c) includes assessing which candidate action is likely to result in the lowest expected growth in regret according to an assumption that the current best observed candidate action is assumed to have zero uncertainty around its mean or expected response performance (See pages 510-11, wherein the candidate action with the expected least regret is represented by zero uncertainty).

Merriman et al. teaches a method and apparatus that considers past performance data when automatically determining the next action to take. Merriman et al. uses a predictive model with which to make a decision, the predictive model using past performance information to deliver optimal actions, thus maximizes utilization of the actions. Eppen et al. discloses the use of regret (or opportunity cost/lost) in the consideration of what action to take with respect to a group of actions based on a set of conditions. It would have been obvious to one of ordinary skill in the art at the time of the invention to use probability distributions and the theory of regret in the iterative predictive model of Merriman et al. in order to increase the efficiency of utilizing advertising/action space by providing a decision framework with which to analyze the various options. See paragraphs 0002, 0008, and 0010 of Merriman et al. and page 503 of Eppen et al.

As per claim 24, Merriman et al. teaches f) applying a temporal depreciation factor to the stored representations of the response performance in order to depreciate the significance of the stored representations over time (See paragraph 0039, wherein a temporal time factor is applied to the actions).

As per claim 25, Merriman et al. wherein step f) includes applying, for each candidate action, a different temporal depreciation factor to the stored representations of the response performance thereof (See paragraph 0039, wherein a factor, such as a seasonal factor, is applied to actions to increase or decrease their relative importance in the problem).

As per claim 30, Merriman et al. discloses wherein the monitored response performance of a respective candidate action in step a) is stored in step b) in a form to enable use of the stored representation of said monitored response performance throughout different components (See paragraphs 0010-1, 0015-6, 0031, 0033-4, wherein the data collected in one component is stored via a server, the server transmitting the data to a component that makes the predictions for the system). However, neither Merriman et al. nor Eppen et al. specifically disclose storing data in a form to enable sharing of the stored representation of said monitored response performance with another system.

Merriman et al. and Eppen et al. are combinable for the reasons set forth above with regards to claim 18.

Further, Merriman et al. discloses storing data in a form that allows the data representing the monitored response performance to be used by different components. Open systems are well known in the Information technology industry and are used to cause interoperability of the system. Therefore, it would have been obvious to one of ordinary skill in the art at the time of

the invention to use open systems architecture in the databases of Merriman et al. in order to increase the usability of the system and the system's data and functions with other system by implementing an interoperable framework.

As per claim 31, Merriman et al. discloses a system having a control apparatus that is programmed to control the objective function of the system according to the method of claim 18 (See figure 1, figure 3, paragraphs 0010-0011, wherein the "advertisement server" controls the system to display certain adds to the user. See also paragraphs 0008, 0017-8, 0033, 0039, 0041-2, wherein an action is chosen based on the current performance and an objective).

As per claim 33, claim 33 recites equivalent limitations to claim 18, and is therefore rejected using the same art and rationale set forth above. Further, Merriman et al. discloses a control apparatus for controlling a system (See figure 1, paragraphs 0010, 0031, wherein a control apparatus is provided).

As per claim 34, Merriman et al. discloses wherein the representation of said monitored response performance contains at least one variable that characterizes the conditions under which the candidate action was performed (See paragraph 0010, 0018-20, 0031, 0033, 0042, which discloses storing monitored feedback concerning the action by the system, wherein the stored response performance considers the variables of context, action type (ie specific ad), etc.).

18. Claims 22 and 26 are rejected under 35 U.S.C. 103(a) as being unpatentable over Merriman et al. (U.S. 2002/0099600) in view of Eppen et al. (*Quantitative Concepts for Management*) and in further view of McClave et al. (*A First Course in Business Statistics*).

As per claim 22, Merriman et al. discloses performing candidate actions (See paragraphs 0010, 0015-6, 0033-4). Merriman et al. further discloses assessing the candidate actions to choose which of the plurality of candidate actions to next performed so as to optimize said objective function by assessing, using a predictive model, which has the current expected best response performance (See paragraphs 0008, 0017-8, 0033, 0039, 0041-2, wherein an action is chosen based on the current known performance). However, Merriman et al. does not expressly disclose assessing which candidate action is likely to result in the lowest expected growth in regret according to an assumption of a Student's distribution and evaluation of Student's t parameters as the basis for estimating probabilities of unequal or equal response states between the candidate action with the current expected best response performance and any other candidate action.

Eppen et al. discloses assessing, using the probability distribution of the response performance of all of said plurality of candidate actions, which candidate action is estimated to result in the lowest expected growth in regret after the chosen candidate action is performed, wherein the actions have unequal response states, based on the probability distribution, between the candidate action with the current expected best response performance and any other candidate action (See page 503, page 504, section 1, 511, section 1, wherein when the decision maker/software knows the probability distribution on the state of nature, regret could be minimized. See also page 512-513). Eppen et al. further discloses the situation where one does not know the specific probability distribution and therefore uses a known probability distribution, such as a minimax criterion, to select a decision that performs the best (See page 511, section 1).

However, Eppen et al. does not expressly disclose using a Student's distribution as the known probability distribution with Student's t parameters as the basis for estimating the probabilities.

McClave et al. discloses determining the distribution of the population using a Student's distribution with Student's t parameters (See pages 297-298).

Merriman et al. teaches a method and apparatus that considers past performance data when automatically determining the next action to take. Merriman et al. uses a predictive model with which to make a decision, the predictive model using past performance information to deliver optimal actions, thus maximizes utilization of the actions. Eppen et al. discloses the use of regret (or opportunity cost/lost) in the consideration of what action to take with respect to a group of actions based on a set of conditions. It would have been obvious to one of ordinary skill in the art at the time of the invention to use probability distributions and the theory of regret in the iterative predictive model of Merriman et al. in order to increase the efficiency of utilizing advertising/action space by providing a decision framework with which to analyze the various options. See paragraphs 0002, 0008, and 0010 of Merriman et al. and page 503 of Eppen et al.

Further, Eppen et al. discloses using a probability distribution associated with the state of nature (i.e. possible outcomes). McClave et al. discloses determining the sample distribution to make reliable decisions using a Student's distribution with t statistics. It would have been obvious to one of ordinary skill in the art at the time of the invention to use a Student's distribution as the distribution in Eppen et al. in order to increase the confidence and reliability of the prediction output by the system, thus decreasing the possibility of opportunity loss. See pages 297-298 of McClave et al. and pages 511-12 of Eppen et al.

As per claim 26, Merriman et al. teaches performing candidate actions (See paragraphs 0010, 0015-6, 0033-4). However, Merriman et al. does not expressly disclose forcing the performance of each candidate action a minimum number of times or at a minimum rate.

Eppen et al. discloses using the probability distribution of the response performance of all of said plurality of candidate actions (See page 503, page 504, section 1, 511, section 1, wherein when the decision maker/software knows the probability distribution on the state of nature, regret could be minimized. See also page 512-513). However, Eppen et al. does not expressly disclose forcing the performance of each candidate action a minimum number of times or at a minimum rate.

McClave et al. discloses determining the sample size needed to make reliable decisions, and thus forcing a sampling of that minimum number of estimates (See pages 316-318, which discusses making a certain number of observations).

Merriman et al. teaches a method and apparatus that considers past performance data when automatically determining the next action to take. Merriman et al. uses a predictive model with which to make a decision, the predictive model using past performance information to deliver optimal actions, thus maximizes utilization of the actions. Eppen et al. discloses the use of regret (or opportunity cost/lost) in the consideration of what action to take with respect to a group of actions based on a set of conditions. It would have been obvious to one of ordinary skill in the art at the time of the invention to use probability distributions and the theory of regret in the iterative predictive model of Merriman et al. in order to increase the efficiency of utilizing advertising/action space by providing a decision framework with which to analyze the various options. See paragraphs 0002, 0008, and 0010 of Merriman et al. and page 503 of Eppen et al.

Further, Eppen et al. discloses using a probability distribution associated with the state of nature (i.e. possible outcomes). McClave et al. discloses determining the sample size needed to make reliable decisions, and thus forcing a sampling of that minimum number of estimates. It would have been obvious to one of ordinary skill in the art at the time of the invention to determine an appropriate sample size of candidate actions and force this number of candidate actions to occur in order to increase the confidence and reliability of the prediction output by the system, thus decreasing the possibility of opportunity loss. See pages 316-318 of McClave et al. and pages 511-12 of Eppen et al.

19. Claim 23 is rejected under 35 U.S.C. 103(a) as being unpatentable over Merriman et al. (U.S. 2002/0099600) in view of Eppen et al. (*Quantitative Concepts for Management*) and in further view of Jameson (U.S. 6,032,123).

As per claim 23, Merriman et al. discloses c) choosing which of the plurality of candidate actions is next performed so as to optimize said objective function by assessing, using a predictive model, empirical data to determine which action will maximize feedback/minimize economic loss after the chosen candidate action is performed based on historical response performances to date (See paragraphs 0008, 0017-8, 0033, 0039, 0041-2, wherein an action is chosen based on the current known performance). However, Merriman et al. does not expressly disclose using a Monte Carlo algorithm to provide understanding of the probability distribution of the response performance of all of the plurality of candidate actions and choosing a candidate action with probability proportional to its contribution to the expected regret estimate.

Eppen et al. discloses using a probability distribution of the response performance of all of said plurality of candidate action to provide an understanding of response performance and choosing a candidate action with its likelihood proportional to its contribution to the regret (See pages 511-513, wherein expected regret of each action is proportional to its contribution to regret. See specifically page 513, section 1). However, Eppen et al. does not expressly disclose using a Monte Carlo algorithm to provide understanding of the probability distribution of the response performance of all of the plurality of candidate actions.

Jameson discloses using a Monte Carlo algorithm to provide understanding using "what if" simulation to facilitate analysis of the response performance of candidate actions (See abstract, column 29, line45-column 30, line 10, wherein Monte Carlo simulation is used on user defined distributions to optimize outputs by simulating potential scenarios).

Merriman et al. teaches a method and apparatus that considers past performance data when automatically determining the next action to take. Merriman et al. uses a predictive model with which to make a decision, the predictive model using past performance information to deliver optimal actions, thus maximizes utilization of the actions. Eppen et al. discloses the use of regret (or opportunity cost/lost) in the consideration of what action to take with respect to a group of actions based on a set of conditions. It would have been obvious to one of ordinary skill in the art at the time of the invention to use probability distributions and the theory of regret in the iterative predictive model of Merriman et al. in order to increase the efficiency of utilizing advertising/action space by providing a decision framework with which to analyze the various options. See paragraphs 0002, 0008, and 0010 of Merriman et al. and page 503 of Eppen et al.

Further, Eppen et al. discloses using a probability distribution associated with the state of nature (i.e. possible outcomes) to predict expected outcomes of regret. Jameson discloses using a Monte Carlo algorithm on user defined distributions to provide understanding of potential outcomes of action using "what if" simulation. It would have been obvious to one of ordinary skill in the art at the time of the invention to determine use Monte Carlo simulation on the defined distribution of Eppen et al. in order to increase the confidence and reliability of the prediction output by understanding better understanding the likelihood of potential outcomes through "what-if" analysis, thus decreasing the possibility of opportunity loss. See abstract of Jameson and pages 511-12 of Eppen et al.

20. Claims 27 and 32 are rejected under 35 U.S.C. 103(a) as being unpatentable over Merriman et al. (U.S. 2002/0099600) in view of Eppen et al. (*Quantitative Concepts for Management*) and in further view of Strickland et al. (U.S. 5,790,407).

As per claims 27 and 32, Merriman et al. and Eppen et al. disclose the method and system, as set forth above in the rejection of claims 18 and 31. Therefore these elements are rejected using the same art and rationale as relied upon above in the rejections of claims 18 and 31. However, neither Merriman et al. nor Eppen et al. disclose the system comprises a robot, the robot controlled according to the method of claim 18.

Strickland discloses control systems for controlling external devices, such as robots, by comparing the response profile of the device to the actual response of the device (See abstract, column 1, lines 15-30, column 3, line 62-column 4, line 15 and lines 20-35).

Merriman et al. and Eppen et al. are combinable for the reasons set forth above with regards to claim 18.

Further, Merriman et al. teaches a method and apparatus that considers past performance data when automatically determining the next action to take. Strickland discloses determining the next action to take when controlling external devices, such as robots, by comparing the response profile of the device to the actual response of the device. It would have been obvious to one of ordinary skill in the art at the time of the invention to use the method for control of Merriman et al. to determine the optimal output for an external device, such as a robot, in order to more accurately produce an optimal output for a device by providing a model with which to analyze the various options. See paragraphs 0002, 0008, and 0010 of Merriman et al.

Conclusion

Applicant's amendment necessitated the new ground(s) of rejection presented in this Office action. Accordingly, **THIS ACTION IS MADE FINAL**. See MPEP § 706.07(a). Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the date of this final action.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Beth Van Doren whose telephone number is (571) 272-6737. The examiner can normally be reached on M-F, 8:30-5:00.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Tariq Hafiz can be reached on (571) 272-6729. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

bvd
October 12, 2006

Beth Van Doren
Beth Van Doren
Patent Examiner
AU 3623